UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/810,196	03/25/2004	Brian Holscher	TRAN-P247	8666

45590 7590 07/07/2010

IV (TRANSMETA)
C/O MURABITO, HAO & BARNES LLP
TWO NORTH MARKET STREET
THIRD FLOOR
SAN JOSE, CA 95113

EXAMINER	
GU, SHAWN X	

ART UNIT	PAPER NUMBER
2189	

MAIL DATE	DELIVERY MODE
07/07/2010	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No.	Applicant(s)	
	10/810,196	HOLSCHER ET AL.	
	Examiner	Art Unit	
	SHAWN X. GU	2189	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE *3* MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *25 May 2010*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-25* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-25* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☒ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☐ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

.... DETAILED ACTION

Response to Amendment

1. This Office action is in response to the claims and remarks filed 25 May 2010. Claims 1-25 are pending. All objections and rejections not repeated below are withdrawn.

Specification

2. Applicant is reminded of the proper content of an abstract of the disclosure.

A patent abstract is a concise statement of the technical disclosure of the patent and should include that which is new in the art to which the invention pertains. If the patent is of a basic nature, the entire technical disclosure may be new in the art, and the abstract should be directed to the entire disclosure. If the patent is in the nature of an improvement in an old apparatus, process, product, or composition, the abstract should include the technical disclosure of the improvement. In certain patents, particularly those for compounds and compositions, wherein the process for making and/or the use thereof are not obvious, the abstract should set forth a process for making and/or use thereof. If the new technical disclosure involves modifications or alternatives, the abstract should mention by way of example the preferred modification or alternative.

The abstract should not refer to purported merits or speculative applications of the invention and should not compare the invention with the prior art.

Where applicable, the abstract should include the following:
(1) if a machine or apparatus, its organization and operation;
(2) if an article, its method of making;
(3) if a chemical compound, its identity and use;
(4) if a mixture, its ingredients;
(5) if a process, the steps.

Extensive mechanical and design details of apparatus should not be given.

3. The abstract of the disclosure is objected to because line 6 refers to purported merits or speculative application of the invention by reciting "to reduce an access latency of the processor". Correction is required. See MPEP § 608.01(b).

Claim Objections

4. Claims 1-8 and 25 are objected to for having the following informalities:

Per claim 1, on line 4, "first" should be inserted between "the memory".

Per claim 25, line 2 appears to contain typographical and/or grammatical errors, and it is suggested that "of" should be deleted and "for" should be replaced by "from".

All dependent claims are objected to for having the same deficiencies contained in the claims they are dependent from.

Claim Rejections - 35 USC § 112

5. The following is a quotation of the first paragraph of 35 U.S.C. 112:

> The specification shall contain a written description of the invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same and shall set forth the best mode contemplated by the inventor of carrying out his invention.

6. Claims 1-17 and 21-25 are rejected under 35 U.S.C. 112, first paragraph, as failing to comply with the written description requirement. The claim(s) contains subject matter which was not described in the specification in such a way as to reasonably convey to one skilled in the relevant art that the inventor(s), at the time the application was filed, had possession of the claimed invention.

Per claim 1, the specification teaches away from the recited limitation "a tracker … recognize processor accesses to a plurality of cache lines within a second memory having a second latency less than the first latency", see application's specification, page 7, lines 1-10, page 8, lines 1-4, page 9, lines 8-12 and 20-26, page 10, lines 2-6, page 11, lines 19-21 and 24-25, page 12, lines 8-9 and 23-27 and page 14, lines 7-13, also see the original claims filed 25 March 2004. In fact, the specification teaches the exact opposite of the claimed invention, by disclosing the step of observing/monitoring CPU accesses cache lines within memory pages stored in the system memory 110, which has the highest latency compared to the L1 and L2 caches. Also see Figures 1 and 3 of the application's drawings, wherein the prefetch unit 120 is shown to be coupled to the system memory and monitoring "cpu memory accesses 301". Also note that although the specification teaches in page 10, lines 2-6 the different means of monitoring bus traffic between the CPU, L1 cache and the L2 cache, this is not adequate support for recognizing processor accesses to a plurality of cache lines in the L1 or L2 cache (the claimed second memory having a second latency less than the first latency). Traffic between CPU and L1/L2 caches can be the result of CPU accesses to cache lines in the system memory. In conclusion, the specification does not teach monitoring/recognizing cache access requests issued by a CPU, instead only CPU requests issued to the system memory are recognized and tracked/monitored.

Per claim 9, the claim is rejected for similar reasons set forth above for claim 1, and the limitation in question is "recognize accesses to cache lines within a second memory having a second latency less than the first latency".

Per claim 22, the means for monitoring, means for using and means for prefetching do not have the written description necessary to support a claim limitation which invokes 35 U.S.C. 112, sixth paragraph as argued by the Applicant in the remarks filed 25 May 2010. In the specification, page 14, lines 7-13, the written description discloses that the prefetcher 920 corresponds to the three means without showing any further physical structure. On the other hand, page 13, lines 10-19 of the specification teaches that the invention can be software based. Furthermore, the drawing in Fig. 9 only shows a box 920 corresponding to the prefetcher. Therefore, there is no adequate written description that discloses physical structure for a prefetcher that corresponds to the monitor means, using means and prefetching means. See MPEP §2181.

All dependent claims are rejected for having the same deficiencies contained in the claims they are dependent from.

7. The following is a quotation of the second paragraph of 35 U.S.C. 112:

> The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

8. Claims 22-24 are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

Per claim 22, the claim is indefinite because the means for monitoring, means for using and means for prefetching do not have the written description necessary to support a claim limitation which invokes 35 U.S.C. 112, sixth paragraph as argued by the Applicant in the remarks filed 25 May 2010. In the specification, page 14, lines 7-

13, the written description discloses that the prefetcher 920 corresponds to the three means without showing any further physical structure. On the other hand, page 13, lines 10-19 of the specification teaches that the invention can be software based. Furthermore, the drawing in Fig. 9 only shows a box 920 corresponding to the prefetcher. Therefore, there is no adequate written description that discloses physical structure for a prefetcher that corresponds to the monitor means, using means and prefetching means. As a result, the metes and bounds of the claim is indefinite as it is unclear what constitutes the physical structure that perform the functions of the monitor means, the using means and the prefetching means. See MPEP §2181.

All dependent claims are rejected as having the same deficiencies as the claims they depend from. Appropriate correction is required.

Claim Rejections - 35 USC § 103

9. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

10. Claims 1-16, 18, 21, 22, 24 and 25 are rejected under 35 U.S.C. 103(a) as being unpatentable over Willke [US 6625696 B1] (hereinafter "Willke"), in further view of Ledebohm et al. [US 7,065,630 B1] (hereinafter "Ledebohm").

Independent Claims:

(A) Per claim 1, Willke teaches an apparatus for a computer system, comprising:

a prefetcher (see col. 2, lines 48-67, storage controller 110 which comprises control logic 112) coupled to a first memory for a requesting device of the computer system, the first memory having a first latency (see col. 2, lines 21-47, Storage Device 120 which can be DRAM or other types of high latency memory for the requesting device 100, see col. 1, lines 20-29 and Fig. 3);

a tracker within the prefetcher and configured to recognize requesting device accesses to a plurality of cache lines within a second memory having a second latency less than the first latency, the second memory operable to supply data to the requesting device responsive to requesting device data requests, wherein the requesting device accesses form a stream type sequential access pattern, and wherein further the tracker is configured to use a bit vector to predictively load a target cache line indicated by the stream-type sequential access pattern from the first memory into the second memory for the requesting device in preparation for the target cache line being requested by the requesting device as part of the stream-type requesting device access pattern (see col. 1, lines 20-42, col. 2, lines 48-67, col. 3, lines 1-37, col. 4, lines 48-51, col. 7, lines 29-40, note that the buffer 114 acts as a cache to temporarily store data fetched from storage device 120, and it contains a set of stored access patterns and prediction and accuracy values for a requesting device 100, this set of patterns and values with the corresponding prefetching function in the storage controller 110 are construed to be a

tracker; see col. 6, lines 41-67, col. 7, lines 1-28 and Fig. 3 for accesses to cache lines; see col. 3, lines 10-37, col. 7, lines 29-40 and col. 8, lines 30-35 for stream-type sequential access pattern; also see Fig. 5, items 510, 520, 540 and 560; also note that buffer 114 has a lower latency than the storage device 120 as consistently taught and/or suggested by the cited reference, because it is a smaller buffer that intercepts data from the storage device 120, and because it is closer to the requesting devices than the storage device 120, Fig. 4 also shows an embodiment wherein buffer 114 is implemented as a cache 414 for the storage device 420).

Willke does not specifically teach the requesting device is a processor except that the requesting device is a peripheral device (see Willke, col. 1, lines 8-9 and lines 46-50 and claim 12). Ledebohm teaches a peripheral device in a computer system wherein the peripheral device is a graphics processor, see Ledebohm, col. 4, lines 4-6 and Fig. 1). Ledebohm's peripheral graphics processor provides a dedicated processing unit for the system's display device and provides functionalities such as generating pixels thereby freeing the main CPU from performing graphics control related tasks and increasing the overall performance of the computer system (see Ledebohm, col. 4, lines 4-60). Therefore, it would have been obvious to one ordinarily skilled in the art a the time of the Applicant's invention to including a graphics processor as one of Willke's peripheral requesting devices in order to provide dedicated display device control functionalities and to improve system performance. As a result of the combined teaching of Willke and Ledebohm, the tracker within the prefetcher would be recognizing stream-type processor accesses and using a bit vector to prefetch

cacheline data from adjacent addresses from a higher latency memory to a lower latency memory in anticipation of processor data requests. Also note that the high latency memory can be Ledebohm's graphics memory 116, Ledebohm's system memory 104, Willke's storage device 120, or a combination of these (see Ledebohm, col. 4, lines 47-67 and col. 6, lines 48-67, the graphics memory 116 is mapped to the virtual memory space).

(B) Per claim 9, Willke teaches an apparatus for a computer system, comprising:

a requesting device (requesting device 100, see col. 1, lines 20-29 and Fig. 3);

a first memory (see col. 2, lines 21-47, Storage Device 120 which can be DRAM or other types of high latency memory) coupled to the requesting device, wherein the first memory having a first latency;

a prefetch unit (see col. 2, lines 48-67, storage controller 110 which comprises control logic 112) coupled to the first memory;

a plurality of trackers included in the prefetch unit, wherein the trackers are respectively configured to recognize requesting device accesses to pages of the first memory (note that Willke teaches DRAM, see col. 2, lines 29-45), and configured to recognize accesses to cache lines within a second memory having a second latency less than the first latency, the second memory operable to supply data to the requesting device responsive to requesting device data requests that form a stream type sequential access pattern; and

the second memory coupled to the prefetch unit, wherein the prefetch unit uses a bit vector to predictively load target cache lines from the first memory into the cache memory to reduce an access latency of the requesting device in preparation for the target cache lines being requested by the requesting device as part the stream-type sequential requesting device access pattern, and wherein the target cache lines are indicated by the stream type sequential access pattern identified by the trackers (see col. 1, lines 20-42, col. 2, lines 48-67, col. 3, lines 1-25, col. 4, lines 48-51, col. 7, lines 29-40, note that the buffer 114 acts as a cache to temporarily store data fetched from storage device 120, and it contains a set of stored access patterns and prediction and accuracy values for each requesting device, this set of patterns and values with the corresponding prefetching function in the storage controller 110 are construed to be a tracker; see col. 6, lines 41-67, col. 7, lines 1-28 and Fig. 3 for accesses to cache lines; see col. 3, lines 10-25, col. 7, lines 29-40 and col. 8, lines 30-35 for stream-type sequential access pattern; also see Fig. 5, items 510, 520, 540 and 560; also note that buffer 114 has a lower latency than the storage device 120 as consistently taught and/or suggested by the cited reference, because it is a smaller buffer that intercepts data from the storage device 120, and because it is closer to the requesting devices than the storage device 120, Fig. 4 also shows an embodiment wherein buffer 114 is implemented as a cache 414 for the storage device 420).

Willke does not specifically teach the requesting device is a processor except that the requesting device is a peripheral device (see Willke, col. 1, lines 8-9 and lines 46-50 and claim 12). Ledebohm teaches a peripheral device in a computer system

wherein the peripheral device is a graphics processor, see Ledebohm, col. 4, lines 4-6 and Fig. 1). Ledebohm's peripheral graphics processor provides a dedicated processing unit for the system's display device and provides functionalities such as generating pixels thereby freeing the main CPU from performing graphics control related tasks and increasing the overall performance of the computer system (see Ledebohm, col. 4, lines 4-60). Therefore, it would have been obvious to one ordinarily skilled in the art a the time of the Applicant's invention to including a graphics processor as one of Willke's peripheral requesting devices in order to provide dedicated display device control functionalities and to improve system performance. As a result of the combined teaching of Willke and Ledebohm, the tracker within the prefetcher would be recognizing stream-type processor accesses and using a bit vector to prefetch cacheline data from adjacent addresses from a higher latency memory to a lower latency memory in anticipation of processor data requests. Also note that the high latency memory can be Ledebohm's graphics memory 116, Ledebohm's system memory 104, Willke's storage device 120, or a combination of these (see Ledebohm, col. 4, lines 47-67 and col. 6, lines 48-67, the graphics memory 116 is mapped to the virtual memory space).

(C) Per claims 18 and 22, it should be clear that the instant claim is already substantially described by claims 1 and 9 as set forth above. Note that Willke's buffer stores multiple stream-type access patterns for multiple requesting devices (see Willke, col. 4, lines 48-57 and col. 7, lines 29-40). Also note that Willke teaches monitoring data transfers between a first memory having a first latency and a second memory coupled to a processor by using a prefetcher, wherein the first memory has a first latency and the second memory has a second latency less than the first latency (see Willke, col. 3, lines 19-37; note that data requested by requesting device 100 is provided to the requesting device 100 while the data is also being stored in buffer 114, and combined this teaching with Fig. 1, it can be shown that the storage controller 110 monitors, or oversees/supervises data transfers between storage device 120 and buffer 114; also note that the data requests by requesting device 100 tracked/monitored by storage controller 110 are essentially data transfers between storage device 120 and buffer 114 because they are segments of the data transfers between requesting device 100 and storage device 120).

Dependent Claims:

(D) Per claims 2 and 10, Willke in view of Ledebohm further teaches each of the trackers include a tag (see Willke, col. 2, lines 48-67, col. 3, lines 1-25, col. 4, lines 48-51, col. 7, lines 29-40; note that any configuration/setting data can be construed as a tag, and there are separate sets of prediction and accuracy values and stored access patterns for each processor) configured to recognize accesses to corresponding cache lines of the first memory by the processor (note that Willke teaches DRAM, see col. 2, lines 29-45, col. 6, line 52-67 and col. 7, lines 1-28).

(E) Per claims 3 and 11, Willke in view of Ledebohm further teaches a plurality of first memory accesses by the processor to the first memory as recognized by the tag are used by the trackers to determine the target cache line for a predictive load into the second memory (see Willke col. 2, lines 48-67, col. 3, lines 1-25, col. 4, lines 48-51, col. 7, lines 29-40 and the rejection of claims 1 and 9 set forth above; the stored access patterns are formed by multiple accesses by the corresponding requesting device; also note that both Willke's storage device 120 and/or Ledebohm's graphics memory 116 can be viewed as a system memory because under the broadest reasonable interpretation a system memory is a memory associated as a system).

(F) Per claims 4, 12 and 24, Willke in view of Ledebohm further teaches consecutive accesses by the processor to adjacent cache lines of a page of the system/first memory are used to determined the target cache line of a stream type access pattern for a predictive load into the second memory, wherein the adjacent cache lines have adjacent addresses, and wherein the target cache line is part of stream-type accesses that formed the stream-type access pattern (see Willke, col. 6,

lines 41-67, col. 7, lines 1-28 and Fig. 3 for accesses to cache lines; see col. 3, lines 10-25, col. 7, lines 29-40 and col. 8, lines 30-35 for stream-type sequential access pattern and adjacent addresses).

(G) Per claim 5, Willke in view of Ledebohm further teaches the first memory comprises a memory block of a plurality of memory blocks of the computer system (see Willke, DRAM and disks, col. 2, lines 29-45, also see col. 6, line 52-67 and col. 7, lines 1-28; a memory block can be a cache line, a page, a byte, or any other unit of memory storage).

(H) Per claims 6 and 13, Willke does not specifically teach the first memory comprises a plurality of 4KB pages but teaches the first memory can be a DRAM (see col. 2, lines 21-40). DRAM page sizes are determined by design choices and system specifications, and page sizes such as 4KB and 8KB are common in the art. Therefore, it would have been obvious to one ordinarily skilled in the art at the time of the Applicant's invention to use 4KB as the page size in Willke's DRAM if dictated by design choice and system specification.

(I) Per claims 7 and 14, Willke in view of Ledebohm further teaches each of the plurality of trackers includes a tag configured to monitor a sub portion of the first memory block/page for accesses by the processor (see Willke, col. 2, lines 48-67, col. 3, lines 1-25, col. 4, lines 48-51, col. 7, lines 29-40; note that any configuration/setting data can be construed as a tag, and there are separate sets of tracking/monitoring, prediction and accuracy values and stored access patterns for each requesting device; each of the stored access patterns only corresponds to a sub portion of the DRAM/the first memory).

(J) Per claim 8, Willke in view of Ledebohm further teaches the first memory is a system memory of the computer system (see Willke, DRAM, col. 2, lines 29-45; also note that both Willke's storage device 120 and/or Ledebohm's graphics memory 116 can be viewed as a system memory because under the broadest reasonable interpretation a system memory is a memory associated as a system).

(K) Per claim 15, Willke further teaches the cache line are 64 byte cache lines (see col. 7, lines 14-15) and a tag is used to monitor half of a page for accesses by the processor (here a page is broadly construed as an unit of storage that is twice the size of a cache line, also note that any configuration/setting data can be construed as a tag, and there are separate sets of prediction and accuracy values and stored access patterns for each requesting device, see the rejection of claim 2 set forth above), but does not teach the cache lines are 128 bytes. However, it is clear that cache line sizes are design dependent and it would have been obvious to one ordinarily skilled in the art

at the time of the Applicant's invention to make Willke's cache like size 128 bytes as a design choice.

(L) Per claim 16, Willke in view of Ledebohm further teaches that the second memory is a prefetch cache memory within the prefetch unit (see Willke, Fig. 1, Buffer 114 is within Storage Controller 110, also see col. 3, lines 19-37).

(M) Per claim 21, Willke in view of Ledebohm further teaches said prefetch unit accesses to first memory are timed to utilize processor-to-system memory idle time (note that all memory accesses timed by a clock such as system clock in a digital computing system; also Willke, see col. 3, lines 10-25 and Fig.1 and 3, processor-to-system memory is idle during the time storage controller 110 accesses storage device 120 for prefetching and satisfying requesting device 100's requests, also storage controller 110 is only able to access storage device 120 when the processor/requesting device 100 is not directly accessing the storage device 120).

(N) Per claim 25, Willke in view of Ledebohm further teaches the prefetcher comprises a prefetch cache operable to be used to load a cache line for the first memory (see Willke, cache 414 in Fig. 4 and buffer 114 in Fig. 1; see col. 3, lines 19-20; also see col. 2, lines 48-67, col. 3, lines 1-25, col. 4, lines 48-51, col. 7, lines 29-40, note that the buffer 114 contains a set of stored access patterns and prediction and accuracy values for prefetching cache lines from a higher latency storage device to a lower latency storage device, hence buffer 114 is "used to" load a cache line from the storage device 120, also see col. 6, lines 41-67, col. 7, lines 1-28 and Fig. 3 for accesses to

cache lines; hence, buffer 114 is used in two different ways to load a cache line from the first memory).

11. Claim 17 is rejected under 35 U.S.C. 103(a) as being unpatentable over Willke in view of Ledebohm, further in view of Microsoft Computer Dictionary (hereinafter "Microsoft").

(O) Per claim 17, Willke in view of Ledebohm does not specifically disclose that the second memory is an L2 cache memory, but teaches that the first is a DRAM as set forth above in the rejection of claim 1. However, Microsoft discloses that a L2 cache is faster than DRAM and Willke's cache memory is a lower latency memory as compared to the higher latency DRAM. Therefore it would have been obvious to one ordinarily skilled in the art at the time of the Applicant's invention to make Willke's second memory a L2 cache because it is faster than DRAM.

12. Claims 19, 20 and 23 are rejected under 35 U.S.C. 103(a) as being unpatentable over Willke in view of Ledebohm, further in view of Brooks [6,081,868] (hereinafter "Brooks").

(P) Per claims 19 and 23, Willke in view of Ledebohm further teaches the computer system includes one main processor (see Willke, Fig. 4, CPU 400) and a peripheral graphics processor (Ledebohm's graphics processor), but does not specifically disclose that there are a plurality of main processors/CPUs and each of the CPUs is coupled to a respective high latency/first memory and a low latency/second

memory. However, Brooks teaches a prefetch system wherein each of a plurality of CPUs is coupled to a respective high latency memory and a low latency memory (see Brooks: Fig 2, a CPU is coupled to a CPU private memory and a CPU cache in each CPU block; CPU Private Memory has higher latency than CPU Cache), in order to provide data storage exclusively for the associated CPU (Brooks: Col 5, Lines 25-30), higher performance and better fault tolerance. Therefore, it would have been obvious to one ordinarily skilled in the art at the time of the Applicant's invention to increase the number of CPUs in Willke's invention and couple each of Willke's plurality of CPUs to a respective high latency/first memory and a low latency/second memory in order to provide increased performance and fault tolerance.

(Q) Per claim 20, the claim recites substantially similar limitations as claims 4, 12 and 24 set forth above.

Response to Arguments

13. Applicant's arguments with respect to claims 1-25 have been fully considered but are not persuasive. For the Applicant's argument regarding claim 1's rejection under 35 U.S.C. 103(a), this Examiner's respectfully points that the limitation "tracker … configured to recognize processor accesses to a plurality of cache lines within a second memory having a second latency less than the first latency" lacks sufficient support from the written description as set forth above in claim 1's rejection under 35 U.S.C 112(1). As a result, claim 1's rejection over Willke in view of Ledebohm is repeated above without giving patentable weight to this particular limitation argued by the Applicant. Claim 9 is treated similarly. As for claims 18 and 22, the claims do not in fact recite the

particular limitation set forth above and argued by the Applicant in the Remarks filed 25 May 2010. Therefore, the Applicant's argument for claim 1 is moot and cannot be applied to claims 18 and 22.

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Shawn Gu whose telephone number is (571) 272-0703. The examiner can normally be reached on 9am-5pm, Monday through Friday.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Reginald Bragdon can be reached on (571) 272-4204. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

/SHAWN X GU/

Shawn X Gu
Patent Examiner
Art Unit 2189

29 June 2010